

SMT-based Induction Methods for Timed Systems

Roland Kindermann, Tommi Junttila, and Ilkka Niemelä

Aalto University

Department of Information and Computer Science

P.O.Box 15400, FI-00076 Aalto, Finland

Roland.Kindermann@aalto.fi, Tommi.Junttila@aalto.fi,

Ilkka.Niemela@aalto.fi

Abstract. Modeling time related aspects is important in many applications of verification methods. For precise results, it is necessary to interpret time as a dense domain, e.g. using timed automata as a formalism, even though the system's resulting infinite state space is challenging for verification methods. Furthermore, fully symbolic treatment of both timing related and non-timing related elements of the state space seems to offer an attractive approach to model checking timed systems with a large amount of non-determinism. This paper presents an SMT-based timed system extension to the IC3 algorithm, a SAT-based novel, highly efficient, complete verification method for untimed systems. Handling of the infinite state spaces of timed system in the extended IC3 algorithm is based on suitably adapting the well-known region abstraction for timed systems. Additionally, k -induction, another symbolic verification method for discrete time systems, is extended in a similar fashion to support timed systems. Both new methods are evaluated and experimentally compared to a booleanization-based verification approach that uses the original discrete time IC3 algorithm.

1 Introduction

In many application areas of model checking, such as analysis of safety instrumented systems, modeling and analyzing in the presence of dense time constructions such as timers and delays is essential. Compared to finite state systems, such timed systems add an extra layer of challenge for model checking tools. In many cases, timed automata [1–3] are a convenient formalism for describing and model checking timed systems. There are many tools, Uppaal [4] to name just one, for timed automata and model checking algorithms for timed automata have been studied extensively during the last two decades, see e.g. [3] for an overview. Most state-of-the-art model checking systems for timed automata use the so-called region abstraction to make a finite state abstraction of the dense time clocks in the automata. These regions are then manipulated symbolically with difference bounded matrices or decision diagram structures (see e.g. [5]).

In this paper our focus is on model checking of safety instrumented systems (see e.g. [6]). Such systems have features that are challenging for the classic timed automata based approach described above. First, safety instrumented systems do typically involve a substantial number of timing related issues. However, such systems are often not best described using automata-like control structures but with a sequential circuit-like control logic. This makes the use of timed automata rather inconvenient in modeling. Second, such systems tend to have a relatively large amount of non-deterministic

input signals which are computationally challenging for model checking tools based on explicit state representation of discrete components (i.e. control location and data).

Hence, we are interested in developing model checking techniques that complement the automata based methods to address these issues. Instead of timed automata, we use a more generic symbolic system description formalism [7] which can be seen as an extension of the classic symbolic transition systems [8] with dense time clock variables and constraints. In our previous work [7], we have experimented with (i) SMT-based bounded model checking (BMC) [9, 10], and (ii) BDD-based model checking based on booleanization of the region abstracted model. These methods were not totally satisfactory as (i) BMC can, in practice, only find bugs, not prove correctness of the system, and (ii) the BDD-based method does not seem to scale well to realistically sized models.

In order to address the computational challenge to develop model checking techniques that can handle timing as well as a substantial amount of non-deterministic input signals and prove correctness, we turn to inductive techniques. The motivation here is the success of temporal induction [11, 12] and, especially, of the IC3 algorithm [13] in the verification of finite state hardware systems. Our approach is to employ SMT solvers instead of SAT solvers as the basic constraint solver technology and apply symbolic region abstraction to handle the dense time clocks in the models. We extend IC3 to timed systems by using linear arithmetics instead of propositional logic and by lifting the concrete states found by the SMT solver to symbolic region level constraints that are further used in the subsequent steps to constrain the search. As a result we obtain a version of IC3 that does not explicitly construct the symbolic region abstracted system but still can exclude whole regions of states at once. We also describe an SMT-based extension of the k -induction algorithm to these kinds of timed systems. In addition, we develop optimizations that allow us to exclude more regions at a time in the SMT-based IC3 algorithm, and to use stronger “simple path” constraints in k -induction.

Our experimental results indicate that SMT-based IC3 can indeed prove much more properties and on much larger models than were possible with our earlier approaches or with SMT-based timed k -induction. Furthermore, when comparing to the approach of using the original propositional IC3 on booleanized region abstracted model, we observe that using richer logics in the SMT framework makes the IC3 algorithm scale much better for timed systems. However, IC3 seems to perform worse than k -induction (and thus BMC) in finding counter-examples to properties that do not hold. This is probably due to its backwards DFS search nature, and leads us to the conclusion of recommending the use of a portfolio approach combining SMT-based BMC and IC3 when model checking these kinds of safety instrumented systems.

2 Symbolic Timed Transition Systems and Regions

We model timed systems with symbolic timed transition systems (STTS) [7], a generic formalism allowing modeling of arbitrary control logic structures, data manipulation, and non-deterministic external inputs. In a nutshell, STTSs can be seen as symbolic transition systems [8] extended with real-valued clocks and associated constraints. By using encoding techniques similar to those in [9, 10], timed automata (and networks of such) can be efficiently translated into STTS [7].

In the following, we use standard concepts of propositional and first-order logics. We assume typed (i.e., sorted) logics, and that formulas are interpreted modulo some background theories (in particular, linear arithmetics over reals); see e.g. [14] and references therein. If $Y = \{y_1, \dots, y_l\}$ is a set of variables and ϕ formula over Y , then $Y' = \{y'_1, \dots, y'_l\}$ is the set of corresponding similarly typed *next-state variables* and ϕ' is obtained from ϕ by replacing each variable y_j with y'_j . Similarly, if ψ is a formula over $Y \cup Y'$, then, for each $i \in \mathbb{N}$, the formula $\psi^{[i]}$ is obtained by replacing y_j with $y_j^{[i]}$ and y'_j with $y_j^{[i+1]}$, of the same types. For example, if $\psi = (c'_2 \leq c_1 + \delta) \wedge x'_1$, then $\psi^{[4]} = (c_2^{[5]} \leq c_1^{[4]} + \delta^{[4]}) \wedge x_1^{[5]}$.

An STTS (or simply a system) is a tuple $\langle X, C, Init, Invar, T, R \rangle$, where

- $X = \{x_1, \dots, x_n\}$ is a finite set of finite domain *state variables*,
- $C = \{c_1, \dots, c_m\}$ is a finite set of real-valued clock variables (or simply *clocks*),
- $Init$ is a formula over X describing the initial states of the system,
- $Invar$ is a formula over $X \cup C$ specifying a state invariant (throughout the paper, we assume the state invariants to be convex, as defined later),
- T is the *transition relation formula* over $X \cup C \cup X'$, and
- R associates each clock $c \in C$ a *reset condition* formula r_c over $X \cup C \cup X'$.

Like in timed automata context, we require that in all the formulas in the system the use of clock variables is restricted to atoms of the form $c \bowtie n$, where $c \in C$ is a clock variable, $\bowtie \in \{<, \leq, =, \geq, >\}$ and $n \in \mathbb{Z}$. Observe that, as in the timed automata context as well, one could use rational constants in systems and then scale them to integers in a behavior and property preserving way. A system is *untimed* if it does not have any clock variables. For the sake of readability only, we do not consider the so-called urgency constraints [7] in this paper.

The semantics of an STTS is defined by its states and how they may evolve to others. A *state* is simply an interpretation over $X \cup C$. A state s is *valid* if it respects the state invariant, i.e. $s \models Invar$. A state s is an *initial state* if it is valid, $s \models Init$, and $s(c) = 0$ for each clock $c \in C$. Given a state s and $\delta \in \mathbb{R}_{\geq 0}$, we denote by $s + \delta$ the state where clocks have increased by δ , i.e. $(s + \delta)(c) = s(c) + \delta$ for each clock $c \in C$ and $(s + \delta)(x) = s(x)$ when $x \in X$. A valid state s may evolve into a successor state u , denoted by $s \longrightarrow u$, if u is also valid and either of the following holds:

1. *Discrete step*: (i) the current and next state interpretations evaluate the transition relation to true, i.e. $\gamma \models T$ where $\gamma(y) = s(y)$ when $y \in X \cup C$ and $\gamma(x') = u(x)$ when $x' \in X'$, and (ii) each clock either resets or keeps its value: for each clock $c \in C$, $u(c) = 0$ if $\gamma \models r_c$ and $u(c) = s(c)$ otherwise.
2. *Time elapse step*: (i) some amount of time elapses: $u = s + \delta$ for some $\delta \in \mathbb{R}_{\geq 0}$, and (ii) the state invariant is respected in the states in between: $s + \mu$ is valid for all $0 < \mu \leq \delta$.

A path is a finite sequence $s_0 s_1 \dots s_l$ of states such that $s_i \longrightarrow s_{i+1}$ holds for each consecutive pair of states in the path. A state is *reachable* if there is a path from an initial state to that state. A *property* P is a formula over the state variables X and the clock variables C , adhering to the same restrictions on the use of clock variables as the

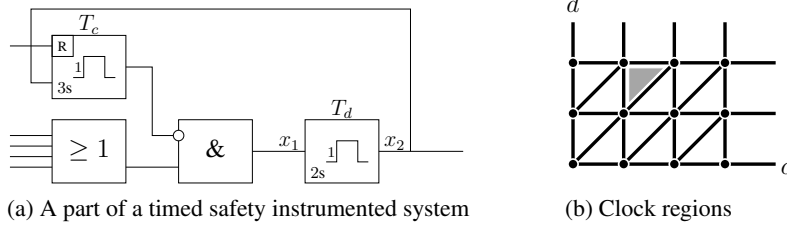


Fig. 1. Illustrations of safety instrumented systems and regions.

system's formulas. In this paper we are interested in solving the problem whether the given state property P is an invariant, i.e. whether P holds in all the reachable states of the system.

As in the contexts of timed automata and linear hybrid automata, we require the state invariants in STTSs to be convex. Formally, a state invariant is *convex* if for all states s and for all $0 \leq \eta \leq \delta$ it holds that whenever $s \models Invar$ and $(s + \delta) \models Invar$, also $(s + \eta) \models Invar$. Thus, a state invariant cannot become false and then true again during a time-elapse step, making condition (ii) of time-elapse steps to always hold. Convexity is easy to test with one call to an SMT solver.

Example 1. As an example, consider an STTS modeling the timer T_d in the safety instrumented system in Fig. 1(a). The STTS has the clock variable d which is reset when a discrete step makes the signal x_1 true, i.e. $r_d = (\neg x_1 \wedge x'_1)$, corresponding to the activation of the timer. The output signal x_2 is initially false, i.e. $Init$ contains the conjunct $(\neg x_2)$. It changes to true when the signal x_1 does and then stays true for two seconds. These properties are captured by the conjunct $(x'_2 \Leftrightarrow (\neg x_1 \wedge x'_1) \vee (x_2 \wedge (d < 2)))$ in the transition relation T . To force the timer output to be reset after two seconds, $Invar$ contains the conjunct $(x_2 \Rightarrow (d \leq 2))$.

Regions. A conceptual tool for handling the infinite state space of an STTS is the region abstraction [1]. For a non-negative real number $a \in \mathbb{R}_{\geq 0}$, let $\text{fract}(a)$ be its fractional part, i.e. $a = \lfloor a \rfloor + \text{fract}(a)$ and $0 \leq \text{fract}(a) < 1$. Let v be an interpretation over C (also called a *clock valuation*). Furthermore, let m_c be the maximum (relevant) value of the clock c , i.e. the largest constant that c is compared to in $Invar$, T , R or P . Two clock valuations v and w belong to the same equivalence class called *region*, denoted by $v \sim w$, if for all clocks $c, d \in C$

1. either (i) $\lfloor v(c) \rfloor = \lfloor w(c) \rfloor$ or (ii) $v(c) > m_c$ and $w(c) > m_c$;
2. if $v(c) \leq m_c$, then $\text{fract}(v(c)) = 0$ iff $\text{fract}(w(c)) = 0$; and
3. if $v(c) \leq m_c$ and $v(d) \leq m_d$, then $\text{fract}(v(c)) \leq \text{fract}(v(d))$ iff $\text{fract}(w(c)) \leq \text{fract}(w(d))$.

Figure 1(b) illustrates the region abstraction for an STTS with two clocks, c and d , with $m_c = 3$ and $m_d = 2$. The thick black lines, thick black dots and the areas in between the thick black lines each represent a different region. The region in which $\lfloor v(c) \rfloor = \lfloor v(d) \rfloor = 1$ and $0 < \text{fract}(v(c)) < \text{fract}(v(d))$ is highlighted in gray.

Two states, s and u , are in the same region, denoted by $s \sim u$, if they agree on the values of the state variables and are in the same region when restricted to clock variables.

Due to the restrictions imposed on the use of clock variables, states in the same region are (i) indistinguishable for predicates, meaning that $u \models \text{Init}$ iff $s \models \text{Init}$, $u \models \text{Invar}$ iff $s \models \text{Invar}$, and $u \models P$ iff $s \models P$ whenever $s \sim u$, and (ii) forward bisimilar: if $s \longrightarrow s'$ and $s \sim u$, then there exists a u' such that $u \longrightarrow u'$ and $s' \sim u'$.

Formula Representation with Combined Steps. To simplify the exposition, to reduce amount of redundancy in paths, and to enable some optimizations, we introduce a *formula representation* for STTSs that exploits a well-known observation: for reachability checking, it is enough to consider paths where discrete steps and time elapse steps alternate, as two consecutive time elapse steps can be merged into one and zero duration time elapse steps can be added in between discrete steps. For a given STTS $\langle X, C, \text{Init}, \text{Invar}, T, R \rangle$, we define the following formulas:

- $\widehat{\text{Invar}} := \text{Invar} \wedge \bigwedge_{c \in C} c \geq 0$. Now $s \models \widehat{\text{Invar}}$ for a state s iff s is a valid state and all clock values are non-negative.
- $\widehat{\text{Init}} := \text{Init} \wedge \bigwedge_{c \in C} c = \hat{c}$ for a free real-valued variable \hat{c} . Now $s \models \widehat{\text{Init}}$ iff, forgetting the state validity requirements, s is a state reachable from an initial state with time elapse steps only.
- $\widehat{T} := T \wedge \delta \geq 0 \wedge \bigwedge_{c \in C} (r_c \Rightarrow c' = \delta) \wedge \bigwedge_{c \in C} (\neg r_c \Rightarrow c' = c + \delta)$. Thus, a state u is reachable from a state s with one discrete step followed by one time elapse step iff $\pi \models \widehat{T}$ for the valuation π on $X \cup C \cup X' \cup C'$ mapping each $z \in X \cup C$ to $s(z)$ and each $z' \in X' \cup C'$ to $u(z')$.

3 k -Induction for Timed Systems

The k -induction method [11, 12] inductively proves a reachability property for a system or discovers a counter-example while trying to prove the property. In the following, we will extend k -induction, which was originally proposed as a verification method for finite-state systems, to a complete verification method for STTS.

As the base case of an inductive proof, k -induction shows that no bad state can be reached within k steps starting from an initial state for some $k \in \mathbb{N}$. As the inductive step, k -induction shows that it is impossible under the transition relation of the system to have a path consisting for k good (property-satisfying) states followed by a bad (property-violating) state. Together, base case and inductive step prove that the property holds in any reachable state.

For an untimed system $\langle X, \emptyset, \text{Init}, \text{Invar}, T, \emptyset \rangle$, both the base case and inductive step can be proven using a SAT solver. The base case holds iff the formula $\text{Init}^{[0]} \wedge \bigwedge_{i=0}^k \text{Invar}^{[i]} \wedge \bigwedge_{i=0}^{k-1} T^{[i]} \wedge \bigwedge_{i=0}^{k-1} P^{[i]} \wedge \neg P^{[k]}$ is unsatisfiable. Likewise, the inductive step holds iff the formula $\bigwedge_{i=0}^k \text{Invar}^{[i]} \wedge \bigwedge_{i=0}^{k-1} T^{[i]} \wedge \bigwedge_{i=0}^{k-1} P^{[i]} \wedge \neg P^{[k]}$ is unsatisfiable. Initially, k -induction attempts an inductive proof with $k = 0$. If unsuccessful, k is increased until the inductive proof succeeds or a counter-example is found while checking the base case. Note that the large overlap both between the formulas for checking

base case and inductive step and between the checks before and after increasing k can be exploited by incremental SAT solvers [12].

While correct, the described approach is not complete due to the fact that the induction step is not guaranteed to hold even if the property checked is satisfied by the system. k -induction can, however, be made complete for finite-state systems by only considering simple (non-looping) paths when checking the inductive step. The most straightforward way to enforce paths to be simple, is to add a quadratic number disequality constraints to the SAT formula, requiring any pair of states to be distinct. Experimental evidence, however, suggests that it is beneficial to only add disequality constraints for pairs of states for which it is observed that disequality constraints are needed [12].

k-induction for STTS. Both base case and inductive step formulas can be applied to an STTS $\langle X, C, \widehat{Init}, \widehat{Invar}, T, R \rangle$ simply by replacing \widehat{Init} , T and \widehat{Invar} in these formulas by \widehat{Init} , \widehat{T} and \widehat{Invar} and using an SMT solver instead of a SAT solver. However, unlike for untimed systems, termination is not even guaranteed when adding disequality constraints. For untimed systems, disequality constraints guarantee termination due to the fact that in a finite state system, there are no simple paths of infinite length and, thus, the simple path inductive step check is guaranteed to be unsatisfiable with sufficiently large k . Timed systems, in contrast, typically have no upper bound for the length of a simple path and, thus, disequality constraints are not sufficient for completeness. However, the infinite state space of an STTS can be split into a finite number of regions. Thus, any reasoning made for finite state systems can be applied to regions of states. In particular, k -induction is complete and correct when only paths that do not visit two states belonging to the same region are considered in the inductive step [15]. By enforcing this property on inductive step paths using region-disequality constraints, complete k -induction can be performed using \widehat{Init} , \widehat{T} and \widehat{Invar} (almost) without modification.

In order to specify that two states of an STTS belong to different regions, region-disequality constraints need to individually constrain the integer and fractional parts of clock values. As only some SMT-solvers, such as Yices [16], allow referring to integer and fractional parts of real-valued variables, we provide a region-disequality constraint encoding that does not rely on such a feature.¹ Instead, we split each clock variable c into two variables: c_{int} represents the integer and c_{frac} the fractional part of c 's value. This “splitting of clocks” requires rewriting of \widehat{Init} , \widehat{T} and \widehat{Invar} by replacing each atom involving a clock with a formula as follows:

Atom	Replacement, $n \in \mathbb{N}$	Atom	Replacement, $n \in \mathbb{N}$
$c < n$	$c_{\text{int}} < n$	$c \leq n$	$c_{\text{int}} < n \vee (c_{\text{int}} = n \wedge c_{\text{frac}} = 0)$
$c > n$	$c_{\text{int}} > n \vee (c_{\text{int}} = n \wedge c_{\text{frac}} > 0)$	$c \geq n$	$c_{\text{int}} > n$
$c = n$	$c_{\text{int}} = n \wedge c_{\text{frac}} = 0$	$c = \hat{c}$	$c_{\text{int}} = \hat{c}_{\text{int}} \wedge c_{\text{frac}} = \hat{c}_{\text{frac}}$
$c' = \delta$	$c_{\text{int}} = \delta_{\text{int}} \wedge c_{\text{frac}} = \delta_{\text{frac}}$		
$c' = c + \delta$	$((c_{\text{frac}} + \delta_{\text{frac}} < 1) \Rightarrow (c'_{\text{int}} = c_{\text{int}} + \delta_{\text{int}} \wedge c'_{\text{frac}} = c_{\text{frac}} + \delta_{\text{frac}})) \wedge$ $((c_{\text{frac}} + \delta_{\text{frac}} \geq 1) \Rightarrow (c'_{\text{int}} = c_{\text{int}} + \delta_{\text{int}} + 1 \wedge c'_{\text{frac}} = c_{\text{frac}} + \delta_{\text{frac}} - 1))$		

¹ In [17] we give an alternative encoding for region-disequality constraints in a BMC setting.

Then, two states with indices i and j can be forced to be in different regions by the following region-disequality constraint $DiffRegion^{[i,j]}$:

$$\begin{aligned} \bigvee_{x \in X} x^{[i]} \neq x^{[j]} \vee \bigvee_{c \in C} (c_{\text{int}}^{[i]} \neq c_{\text{int}}^{[j]} \wedge (\neg max_c^{[i]} \vee \neg max_c^{[j]})) \\ \vee \bigvee_{c \in C} (\neg max_c^{[i]} \wedge \neg (c_{\text{fract}}^{[i]} = 0 \Leftrightarrow c_{\text{fract}}^{[j]} = 0)) \\ \vee \bigvee_{c \in C} \bigvee_{d \in C \setminus \{c\}} (\neg max_c^{[i]} \wedge \neg max_d^{[j]} \wedge \neg ((c_{\text{fract}}^{[i]} \leq d_{\text{fract}}^{[i]}) \Leftrightarrow (c_{\text{fract}}^{[j]} \leq d_{\text{fract}}^{[j]}))) \end{aligned}$$

where the shorthand $max_c^{[i]} := c_{\text{int}}^{[i]} > m_c \vee (c_{\text{int}}^{[i]} = m_c \wedge c_{\text{fract}} > 0)$ detects whether the clock c exceeds its maximum relevant value m_c .

4 IC3 for Timed Systems

In this section, we first describe the IC3 algorithm [13] for untimed finite state systems (see also [18] for an alternative, complementary account of the algorithm). We then show how it can be extended for verifying timed systems by using region abstraction and SMT solvers.

Like k -induction, the IC3 algorithm tries to generate an inductive proof for a given state property P on an untimed system $S = \langle X, \emptyset, Init, Invar, T, \emptyset \rangle$. But unlike the unrolling-based approach used by k -induction, proofs generated by the IC3 algorithm only consists of a single formula *Proof* satisfying three properties: (a) *Proof* is satisfied by any initial state of S , (b) *Proof* is satisfied by any successor of any state satisfying *Proof*, and (c) *Proof* $\Rightarrow P$. Properties (a) and (b) serve as base case and inductive step for showing that the set of states satisfying *Proof* is an over-approximation of the states reachable in S while property (c) proves that any reachable state satisfies P .

In order to generate a proof, the IC3 algorithm builds a sequence of sets of formulas $F_0 \dots F_k$ satisfying certain properties. Eventually, one of these sets becomes the proof *Proof*. Each F -set represents the set of states satisfying all its formulas. The properties satisfied by the sequence are (i) $Init \wedge Invar \Rightarrow F_0$, (ii) $F_i \Rightarrow F_{i+1}$, (iii) $F_i \Rightarrow P$, and (iv) $F_i \wedge Invar \wedge T \wedge Invar' \Rightarrow F_{i+1}'$. The basic strategy employed by the IC3 algorithm is to add clauses to the F_i -sets in a fashion that keeps properties (i) to (iv) intact until $F_k \wedge Invar \wedge T \wedge Invar' \Rightarrow P'$. In this situation, k can be increase by appending $\{P\}$ to the sequence. The algorithm terminates once $F_i = F_{i+1}$ for some i and provides F_i as a proof. Upon termination, properties (i) and (ii) imply proof-property (a), property (iv) and the termination condition $F_i = F_{i+1}$ imply property (b) and property (iii) implies property (c). Note that, in practice, property (ii) is enforced by adding any formula added to a given F -set also to all F -sets with lower index, i.e. $F_i \subseteq F_{i-1}$.

After sketching the basic strategy, we will now take a closer look at the algorithm. Note, however, that the description given is only a simplified version of the algorithm that focuses on the aspects that are relevant with respect to extending it for STTS. Figure 2 shows the main loop of the IC3 algorithm. In each iteration, the algorithm first checks whether or not it is currently possible to extend the sequence of F -sets by appending P . Note that as appending P will never result in properties (i) to (iii)

```

1: loop
2:   if  $F_k \wedge Invar \wedge T \wedge Invar' \wedge \neg P'$  is UNSAT then
3:      $k := k + 1$ 
4:     add  $F_k \leftarrow \{P\}$  to sequence of  $F$ -sets
5:     propagate()
6:     if  $F_i = F_{i+1}$  for some  $i$  then
7:       return true {Property holds}
8:   else
9:      $s \leftarrow$  predecessor of a bad state extracted from the model
10:     $success \leftarrow blockState(s)$ 
11:    if  $\neg success$  then
12:      return false {Property violated}

```

Fig. 2. The main loop of IC3

being violated, it is sufficient to check whether extending the sequence would violate property (iv). A corresponding SAT call can be found in Line 2 of Fig. 2. If the SAT call indicates that the sequence can safely be extended, the sequence is extended in Lines 3 and 4. In the next step, the F -set sequence, clauses may be propagated from F -sets to subsequent sets in the sequence. While this step is vital for termination, a more detailed description is omitted here for space limitations. After propagation, the algorithm's termination condition is checked in Line 6.

Of course, the SAT check in Line 2 may as well indicate that the F -sequence may currently not be extended without violating property (iv). In this case, a state s that satisfies F_k and has a bad successor can be extracted from the model returned by the SAT solver. As s prevents the sequence from being extended, the algorithm attempts to drop s from (the set of states represented by) F_k by adding a clause that implies $\neg s$ ². The corresponding subroutine call, $blockState(s)$, may also need to add further clauses also to other F -sets than F_k in order to ensure that properties of the sequence remain satisfied.

The $blockState(s)$ subroutine, outlined in Figure 3, operates on a list of proof obligations, each being a pair of a state and an index. An obligation $\langle s, i \rangle$ indicates that it is necessary to drop s from F_i before the main loop of the algorithm can continue. Initially, the only proof obligation is to drop the state provided as an argument from F_k . For any proof obligation $\langle s, i \rangle$, the $blockState$ subroutine in Line 6 checks whether or not s has a predecessor z in F_{i-1} . Such a predecessor prevents s from being excluded from F_i without violating property (iv). Thus, if a predecessor is found, the obligation $\langle s, i \rangle$ can not be fulfilled immediately and is added to the set of open obligations again in Line 8. Furthermore, z has to be excluded from F_{i-1} before s can be excluded from F_i . This is reflected by the obligation $\langle z, i - 1 \rangle$ also being added to the set of open obligations in Line 9.

If the SAT call in Line 6 is unsatisfiable, then s has no predecessor in F_{i-1} and can safely be excluded from F_i without violating property (iv). The state s is excluded by

² In a slight abuse of notation, we interpret a state s as formula $\bigwedge_{y \in C \cup X} y = s(y)$ where appropriate.


```

1:  $Q \leftarrow$  priority queue containing  $\langle s, k \rangle$ 
2: while  $Q$  not empty do
3:    $s, i \leftarrow Q.popMin()$ 
4:   if  $i = 0$  then
5:     return false {Counter-example found}
6:   if  $F_{i-1} \wedge \neg s \wedge T \wedge Invar \wedge Invar' \wedge s'$  is SAT then
7:      $z \leftarrow$  predecessor of  $s$  extracted from the model
8:      $Q.add(\langle s, i \rangle)$ 
9:      $Q.add(\langle z, i - 1 \rangle)$ 
10:  else
11:     $F_i.add(generalize(\neg s))$ 
12:    if  $i < k$  then
13:       $Q.add(\langle s, i + 1 \rangle)$ 
14: return true

```

Fig. 3. The $blockState(s)$ sub-routine

adding a generalization of the clause $\neg s$ to F_i . More precisely, the algorithm attempts to drop literals from $\neg s$ in a way that preserves properties (i) and (iv) before adding the resulting clause to F_i . Without this generalization step, states would be excluded one at a time from the F -sets resulting in a method akin to explicit state model checking.

So far, it has been assumed that P holds. If this is not the case, then the main loop will eventually pass a predecessor of a bad state reachable in S to $blockState$. In such a situation, $blockState$ essentially performs a backwards depth-first search that eventually leads to an initial state of S , which is detected in Line 4. Note that it is straightforward to extract a counter-example from the proof obligations if it is detected that P does not hold.

Note that, while sufficient for explain our extensions, only a simplified version of the IC3 algorithm has been described. Most notably, the complete version of the algorithm additionally aims to satisfy proof obligations for multiple successive F -sets at a time if possible and performs generalization based on unsatisfiable cores obtained from SAT calls in various locations. For a description of these techniques as well as complete arguments for correctness and completeness of the approach refer to [13, 18].

Extending IC3 for timed systems. As was the case with k -induction, the key to extending the IC3 algorithm to timed systems is the region abstraction. Again, we will use an SMT-solver instead of a SAT-solver and the combined step encoding \widehat{Init} , \widehat{Invar} , and \widehat{T} on an STTS will replace $Init$, $Invar$ and T . To operate on the region level, we lift each concrete state in a satisfying interpretation returned by the SMT solver into to the region level in the IC3 algorithm code whenever it is passed back to the SMT solver again. To do this, given a state s , we construct a conjunction \tilde{s} of atoms such that \tilde{s} represents all the states in the same region as s , i.e. for any state u it holds $u \models \tilde{s}$ iff $u \sim s$. Formally, \tilde{s} is the conjunction of the atoms given by the following rules:

1. For each state variable $x \in X$, add the atom $(x = s(x))$.
2. For each clock c with $s(c) > m_c$, add the atom $(c > m_c)$.

3. For each clock c with $s(c) \leq m_c$ and $\text{fract}(s(c)) = 0$, add the atoms $(c \leq s(c))$ and $(c \geq s(c))$. Two atoms are added instead of $(c = s(c))$ so that the clause generalization sub-routine has more possibilities for relaxing $\neg\tilde{s}$.
4. For each clock c with $s(c) < m_c$ and $\text{fract}(s(c)) \neq 0$, add the atoms $(c > \lfloor s(c) \rfloor)$ and $(c < \lceil s(c) \rceil)$.
5. For each pair c, d of distinct clocks with $s(c) \leq m_c$, and $s(d) \leq m_d$,
 - (a) if $\text{fract}(s(c)) = \text{fract}(s(d))$, add the atoms $(d \leq c - \lfloor s(c) \rfloor + \lfloor s(d) \rfloor)$ and $(d \geq c - \lfloor s(c) \rfloor + \lfloor s(d) \rfloor)$, and
 - (b) if $\text{fract}(s(c)) < \text{fract}(s(d))$, add the atom $(d > c - \lfloor s(c) \rfloor + \lfloor s(d) \rfloor)$.

What is especially convenient here is that, unlike in the region-disequality constraints in k -induction, there is no need to directly access the integral and fractional parts of clock variables in \tilde{s} because \tilde{s} considers one fixed region. Indeed, all the atoms concerning clock variables will fall in the difference logic fragment of linear arithmetics over reals, having very efficient decision procedures available [19, 20].

We now let the IC3 algorithm operate as in the untimed case except the satisfiability calls are changed to operate on the region level. Especially, the formula in Line 6 of Fig. 3 is modified to $F_{i-1} \wedge \neg\tilde{s} \wedge T \wedge \text{Invar} \wedge \text{Invar}' \wedge \tilde{s}'$ so that it operates on the region level, trying to find a predecessor state in the F_{i-1} -set for *any* state in the same region as s . Furthermore, in Line 11 the clause generalization is called with the clause $\neg\tilde{s}$ that represents all the states that are in a different region than s ; thus we exclude at least all states in the region of s from F_i . No major modifications are required in the clause generalization mechanisms but they can handle clock atoms in the same way the state variable literals are handled.

Example 2. Consider the STTS for the system in Fig. 1(a) discussed in Ex. 1. For the state $s = \{x_1 \mapsto \text{false}, x_2 \mapsto \text{true}, c \mapsto 1.4, d \mapsto 1.65, \dots\}$ we obtain the conjunction $\tilde{s} = (\neg x_1 \wedge x_2 \wedge (c > 1) \wedge (c < 2) \wedge (d > 1) \wedge (d < 2) \wedge (d > c) \wedge \dots)$ that represents all the states in the region of s . The clause that excludes the whole region of s is simply $(x_1 \vee \neg x_2 \vee (c \leq 1) \vee (c \geq 2) \vee (d \leq 1) \vee (d \geq 2) \vee (d \leq c) \vee \dots)$.

The soundness of the timed IC3 algorithm can be argued as follows. We say that a formula ϕ over $X \cup C$ *respects regions* if for all states s and u , $s \sim u$ implies that $\phi \models s$ iff $\phi \models u$. By construction, a state property P as well as \tilde{s} and $\neg\tilde{s}$ for any state s all respect regions. Furthermore, any sub-clause of $\neg\tilde{s}$ returned by the clause generalization sub-routine also respects regions. As a result all the clauses in the F -sets respect regions and thus the F -sets exclude whole regions only. Furthermore, the modified formula $F_{i-1} \wedge \neg\tilde{s} \wedge T \wedge \text{Invar} \wedge \text{Invar}' \wedge \tilde{s}'$ in Line 6 of Fig. 3 is unsatisfiable iff no state in the same region as s can be reached from the F_{i-1} -set; thus excluding the whole region in Line 11 is correct.

Because the number of regions is finite, only a finite number of clauses can be added to any F -set. As a result, the argument for termination given in [18] can be applied to the timed IC3 algorithm as well.

5 Optimizations by Excluding Multiple Regions

We now describe optimizations for timed IC3 and k -induction that sometimes al-

low us to exclude more regions at once during the clause generalization and in the region-disequality constraints, respectively. They both exclude time-predecessor regions of a region, i.e. regions from whose states one can reach the other region by just letting time pass. As an example, all the light regions and the dark gray region (with $c = 3$ and $d > 2$) in Fig. 4 are time-predecessors of the dark gray region. Formally, we define that a clock valuation w is in a time-predecessor region of the clock valuation v , denoted by $w \preceq v$, if for all clocks $c, d \in C$ all the following hold:

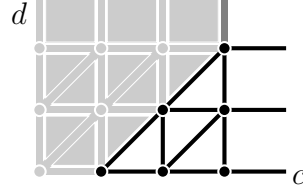


Fig. 4: A time-predecessor clock region

1. Either (i) $v(c) > m_c$, (ii) $\text{fract}(v(c)) = 0$ and $w(c) \leq v(c)$, or (iii) $\text{fract}(v(c)) > 0$ and $w(c) < \lceil v(c) \rceil$.
2. If $v(c) \leq m_c$ and $v(d) \leq m_d$, then
 - (i) $\text{fract}(v(c)) = \text{fract}(v(d))$ implies $w(d) = w(c) - \lfloor v(c) \rfloor + \lfloor v(d) \rfloor$, and
 - (ii) $\text{fract}(v(c)) < \text{fract}(v(d))$ implies $w(d) > w(c) - \lfloor v(c) \rfloor + \lfloor v(d) \rfloor$ and $w(d) < w(c) - \lfloor v(c) \rfloor + \lfloor v(d) \rfloor + 1$.
3. If $v(c) \leq m_c$ and $v(d) > m_d$, then
 - (i) $\text{fract}(v(c)) = 0$ implies $w(d) > w(c) - \lfloor v(c) \rfloor + m_d$, and
 - (ii) $\text{fract}(v(c)) > 0$ implies $w(d) > w(c) - \lfloor v(c) \rfloor + m_d - 1$.

Observe that \preceq is a reflexive relation. A state u is in a time-predecessor region of another state s , denoted by $u \preceq s$, if they agree on the values of the state variables and, when restricted to the clock variables, u is in a time-predecessor region of s .

Application to IC3. The timed variant of the IC3 algorithm described in a previous section excludes an entire region from an F -set once a state inside that region (and thus the whole region) has been found to be unreachable from the previous F -set. In this section, we will argue that it is actually possible to exclude all the time-predecessor regions at the same time. By excluding more than one region, the F -sets potentially shrink faster which can lead to improved execution times. This optimization to the IC3 algorithm is based on the following lemma:

Lemma 1. *Let s be a valid state. If none of the states in the region of s can be reached from an initial state with one time elapse step followed by n combined steps, then none of the valid states in the time-predecessor regions of s can, either.*

Proof. Assume that a valid state $p \preceq s$ is reachable in that way. Thus, (i) p satisfies *Invar*, (ii) there is a $\delta \in \mathbb{R}_{\geq 0}$ such that $p + \delta \sim s$ as $p \preceq s$, (iii) $p + \delta$ satisfies *Invar* as s does, and (iv) all the states “in between” p and s (i.e. all the states $p + \delta'$ with $0 < \delta' < \delta$) satisfy the convex *Invar*, too. Therefore, $p + \delta$ is reachable from an initial state with one time elapse step followed by n combined steps by just “extending” the last time elapse step by δ units. This gives a contradiction as $p + \delta \sim s$. \square

Any state considered by the IC3 algorithm is extracted from a model of a SMT formula containing the system’s state invariant as a conjunct and hence satisfies the invariant. In addition, the \hat{Init} and \hat{T} formulas used in timed IC3 capture initial states followed by one time elapse step and combined steps, respectively. Thus, Lemma 1 is applicable to any state found unreachable by the IC3 algorithm and justifies the dropping

of all the time-predecessor regions at the same time. Given a state s , we can construct a conjunction \tilde{s}_{\prec} of atoms such that \tilde{s}_{\prec} represents all the states in the time-predecessor regions of s . Formally, \tilde{s}_{\prec} is obtained by instantiating the definition of \prec for a concrete state s and is the conjunction of the atoms given by the following rules:

1. For each state variable $x \in X$, add the atom $(x = s(x))$.
2. For each clock c with $s(c) \leq m_c$ and $\text{fract}(s(c)) = 0$, add the atom $(c \leq s(c))$.
3. For each clock c with $s(c) < m_c$ and $\text{fract}(s(c)) \neq 0$, add the atom $(c < \lceil s(c) \rceil)$.
4. For each pair c, d of distinct clocks with $s(c) \leq m_c$ and $s(d) \leq m_d$,
 - (a) if $\text{fract}(s(c)) = \text{fract}(s(d))$, add the atoms $(d \leq c - \lfloor s(c) \rfloor + \lfloor s(d) \rfloor)$ and $(d \geq c - \lfloor s(c) \rfloor + \lfloor s(d) \rfloor)$, again using two literals to encode equality for additional clause relaxation possibilities, and
 - (b) if $\text{fract}(s(c)) < \text{fract}(s(d))$, add the atoms $(d > c - \lfloor s(c) \rfloor + \lfloor s(d) \rfloor)$ and $(d < c - \lfloor s(c) \rfloor + \lfloor s(d) \rfloor + 1)$.
5. For each pair c, d of distinct clocks with $s(c) \leq m_c$ and $s(d) > m_d$,
 - (a) if $\text{fract}(s(c)) = 0$, add the atom $(d > c - \lfloor s(c) \rfloor + m_d)$; and
 - (b) if $\text{fract}(s(c)) > 0$, add the atom $(d > c - \lfloor s(c) \rfloor + m_d - 1)$.

Now \tilde{s}_{\prec} and $\neg\tilde{s}_{\prec}$ can be used instead of s and $\neg s$ in SMT calls and as argument for clause generalization. Observe that \tilde{s}_{\prec} is also in the difference logic fraction of linear arithmetics and does not need to refer to integral or fractional parts of clocks.

Example 3. Consider again the STTS for the system in Fig. 1(a) discussed in Ex. 1. For the state $s = \{x_1 \mapsto \text{false}, x_2 \mapsto \text{true}, c \mapsto 3.0, d \mapsto 2.7, \dots\}$ in the dark gray clock region in Fig. 4, we get the conjunction $\tilde{s} = (\neg x_1 \wedge x_2 \wedge (c \leq 3) \wedge (d > c - 1) \wedge \dots)$ representing all the states in the time-predecessor regions.

Application to k -induction. The idea of excluding time-predecessor regions can also be applied to k -induction. This is based on the following lemma, stating that a path of combined steps can be compressed into a shorter region-equivalent one if a state in it is in the time-predecessor region of a later state:

Lemma 2. *Let $s_0 s_1^d \dots s_{i-1}^d s_i^d s_{i+1}^d \dots s_j^d s_j \dots s_k^d$ be a path such that (i) s_0 is an initial state, (ii) $s_i^d \prec s_j^d$, (iii) each step between s_l and s_{l+1}^d is a time elapse step, and (iv) each step between s_l^d and s_l is a discrete step. Then $s_0 s_1^d \dots s_{i-1}^d s_{i-1} u_j^d u_j \dots u_k^d$ with $u_j^d \sim s_j^d$ for all $j \leq l \leq k$ and $u_j \sim s_j$ for all $j \leq l < k$ is also a path.*

Proof. As $s_i^d \prec s_j^d$ and the state invariants are convex, the time elapse step from s_{i-1} to s_i^d can be “extended” so that a state $s_i^d + \delta \sim s_j^d$ is reached instead. Letting u_j^d equal $s_i^d + \delta$, the existence of the requested postfix $u_j^d u_j \dots u_k^d$ of the path follows from the forward bisimilarity of the states in the same region. \square

Now this implies that in timed k -induction we can use, instead of the region-disequality formula $\text{DiffRegion}^{[i,j]}$, a stronger formula $\text{DiffRegion}_{\prec}^{[i,j]}$ excluding the state $s^{[i]}$ from being in a time-predecessor region of the state $s^{[j]}$ when $i < j$. We omit the details but this formula can be obtained from the definition of the \prec relation in a similar way as the $\text{DiffRegion}^{[i,j]}$ formula was obtained from the definition of \sim in Sect. 3.

Table 1. Verification times in seconds for industrial benchmarks. Blank cells indicate that the respective property can not be verified on the respective size model.

Property	Satisfied	Full					Medium size submodel					Small submodel				
		Timed IC3, extended r.	Timed IC3, basic regions	k-induction, basic regions	k-induction, extended r.	Booleaniza- tion + IC3	Timed IC3, extended r.	Timed IC3, basic regions	k-induction, basic regions	k-induction, extended r.	Booleaniza- tion + IC3	Timed IC3, extended r.	Timed IC3, basic regions	k-induction, basic regions	k-induction, extended r.	Booleaniza- tion + IC3
1	yes	0.55	0.55	0.5	0.53	181.95	0.32	0.33	0.42	0.29	16.43	0.21	0.21	0.29	0.19	8.51
2	yes	0.55	0.51	0.47	0.5	timeout	0.3	0.32	0.33	0.34	timeout	0.21	0.21	0.23	0.22	1459.6
3	yes	0.57	0.56	0.49	0.49	timeout	0.32	0.33	0.36	0.36	timeout	0.22	0.21	0.23	0.24	timeout
4	yes	0.54	0.57	0.48	0.58	timeout	0.33	0.34	0.35	0.34	timeout					
5	yes	0.68	0.55	0.62	0.54	191.81	0.32	0.31	0.3	0.31	18.1					
6	yes	0.57	0.58	0.6	0.5	timeout	0.34	0.33	0.37	0.38	timeout					
7	yes	0.57	0.6	0.51	0.55	timeout										
8	no	timeout	timeout	2.21	2.24	timeout	0.3	0.29	0.35	0.33	timeout	0.21	0.2	0.33	0.23	2367.16
9	no	0.62	0.56	0.7	0.65	194.25	0.32	0.31	0.29	0.27	4.47	0.21	0.23	0.29	0.2	2.05
10	no	timeout	timeout	2.36	2.15	165.16	0.31	0.32	0.41	0.31	17.07	0.2	0.23	0.21	0.2	8.58
11	no	0.53	0.54	0.62	0.65	169.91	0.33	0.35	0.41	0.32	17.53					
12	no	timeout	timeout	2.21	2.11	timeout	0.32	0.31	0.34	0.33	timeout					
13	no	timeout	timeout	2.24	2.17	timeout	0.32	0.31	0.46	0.35	timeout					
14	no	0.57	0.57	0.63	0.65	170.87	0.32	0.36	0.42	0.33	17.8					
15	no	0.56	0.59	0.85	0.66	81.07										

6 Experiments

To determine the usefulness of the described methods, they were evaluated experimentally. Specifically, we were interested in the following questions: how do the methods perform and scale (i) in the area they were designed for, i.e. timed systems with a large amount of non-determinism; (ii) compared to each other; (iii) compared to using discrete time verification methods in a semantics-preserving way; and (iv) outside the area they were designed for, i.e. on models with a low amount of non-determinism?

Setup. Timed k -induction and the timed IC3 algorithm were implemented in Python, each supporting both region encoding variants. Using a more efficient programming language like C is likely to yield only moderate execution time improvements due to a significant fraction of the time being spent by the SMT solver. As an SMT-solver, Yices [16] version 1.0.31 was used. All experiments were executed on Linux computers with AMD Opteron 2435 CPUs limited to one hour of CPU time and 2 GB of RAM.

Industrial benchmark. The first benchmark used is a model of an emergency diesel generator intended for the use in a nuclear power plant. The full model and two sub-models, which are sufficient for the verification some of the properties, were used. The numbers of clocks and state variables are 24 and 130 for the full model, 7 and 64 for the first and 6 and 36 for the second sub-model. The industrial model has been studied previously and found very challenging. Only some partial results [21] have been obtained using the model checker NuSMV [22] by abstracting model based on its component structure and then using a discrete time version of the model. Efforts to verify the abstracted model using the real time model checker Uppaal [4] were even less successful [21]. Likewise, a booleanization-based attempt to verify the smallest sub-model was unable to verify all properties [7].

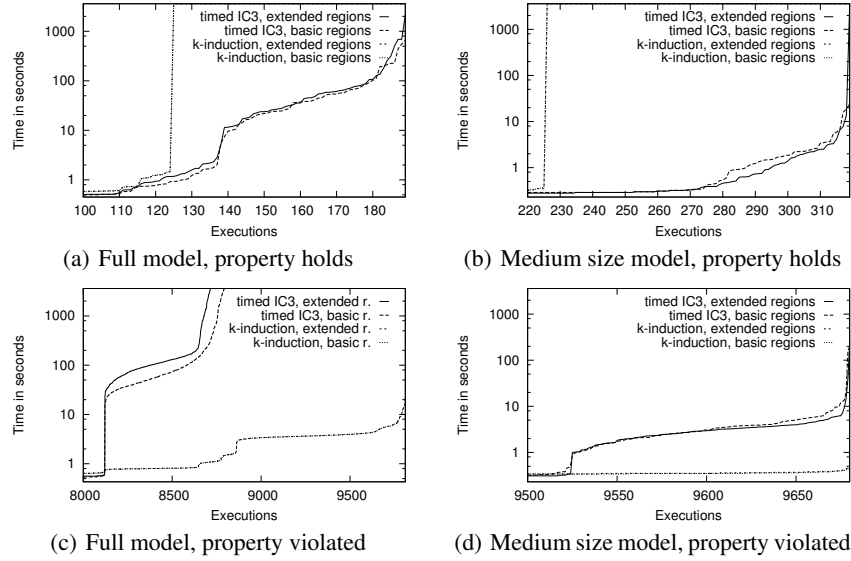


Fig. 5. Time required to verify by numbers of properties for randomly generated properties

All four variants of the methods introduced in this paper were applied to the industrial model. Additionally, the original IC3 implementation [13] in combination with a semantics-preserving booleanization approach [7] was used. Table 1 shows the resulting execution times. k -induction did not exceed three seconds for any property. The timed IC3-approaches performed similarly for most properties but timed out four times. Both real-time verification methods performed significantly better than the booleanization / IC3 combination, illustrating that development of specialized real time verification methods is worthwhile.

Random properties. While the industrial benchmark showed that the methods work well in the area they were designed for, execution times were generally too low to compare the different methods and variants. Therefore, 10000 additional random properties were generated each for the full model and the medium size sub-model, each property being a three literal clause using state variables and / or clocks. Figure 5 shows the resulting execution times. Note that all methods timed out for one property on the medium size model, which then could not be considered in the plots due to not being known whether it holds. For violated random properties, k -induction performed very well, due to its bounded model checking component. For properties that hold, in contrast, timed IC3 performed significantly better. Executing both methods (or timed IC3 and bounded model checking) in parallel could combine their strengths.

Using time-predecessor regions made no difference for k -induction. For the timed IC3 algorithm, their effect depended on the size of the model used. A performance increase was observed for the medium size model, contrasting a performance decrease for the large model. A likely explanation for this behavior is the large number of clocks

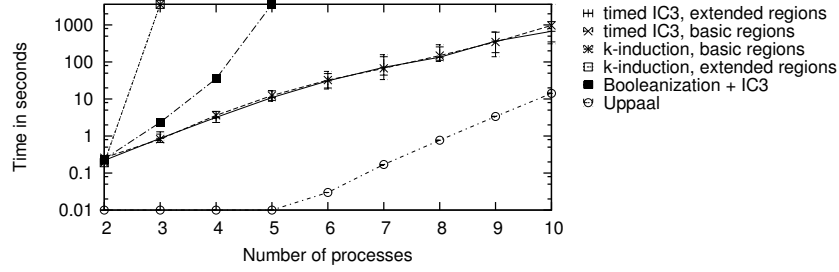


Fig. 6. Verification time the Fischer protocol (min, max and median of 11 executions)

used in the large model. While the time-predecessor region encoding uses fewer literals referring to a single clock than the original region encoding, it contains more literals comparing two clocks. Thus, the size of clauses grows quicker in the number of clocks used for time-predecessor regions, which eventually outweighs the gain of excluding more states at once.

Fischer protocol. As a third benchmark, the Fischer mutual exclusion protocol, a standard benchmark for timed verification, was used. In addition to the five methods used for the industrial method, Uppaal [4] version 4.0.11, a model checker for networks of timed automata, was used. Unlike the industrial benchmark, the Fischer protocol is fairly deterministic and, thus, could be expected to favor Uppaal over the fully-symbolic methods. Figure 6 shows the execution times for verifying the Fischer property with a varying number of processes. While timed IC3 was, unsurprisingly, significantly slower than Uppaal, it scaled similarly, i.e. the runtime increased at a similar rate. k -induction timed out at three processes already while the booleanization-based approach showed exponential runtime growth and timed out at five processes.

7 Conclusion

This paper introduces two verification methods for symbolic timed transition systems: a timed variant of the IC3 algorithm and an adapted version of k -induction. Furthermore, a potential optimization to both methods is devised.

Both methods were able to verify properties on an industrial model verification of which had been found in previous attempts intractable and outperformed a booleanization-based approach significantly. Random properties on the same model revealed that the timed IC3 variant performs better for satisfied properties while timed k -induction performs better on violated properties. The experiments suggest that executing timed IC3 in parallel with bounded model checking would yield excellent performance for the verification of large, non-deterministic real-time systems.

Additionally, the proposed methods were evaluated on another family of benchmark, the Fischer mutual exclusion protocol with a varying number of processes. This family has only a small amount of non-determinism and the runtime of the methods was higher than that of the timed automata model checker Uppaal. However, the timed IC3 algorithm was found to have similar good scaling as Uppaal.

References

1. Alur, R., Dill, D.L.: A theory of timed automata. *Theor. Comput. Sci.* **126**(2) (1994) 183–235
2. Alur, R.: Timed automata. In: *Proc. CAV 1999*. Volume 1633 of LNCS., Springer (1999) 8–22
3. Bengtsson, J., Yi, W.: Timed automata: Semantics, algorithms and tools. In: *Lectures on Concurrency and Petri Nets*. Volume 3098 of LNCS., Springer (2003) 87–124
4. Behrmann, G., David, A., Larsen, K.G.: A tutorial on Uppaal. In: *Proc. FM-RT 2004*. Volume 3185 of LNCS., Springer (2004) 200–236
5. Wang, F.: Efficient verification of timed automata with BDD-like data structures. *International Journal on Software Tools for Technology Transfer* **6**(1) (2004) 77–97
6. Gruhn, P., Cheddie, H.L.: *Safety Instrumented Systems: Design, Analysis, and Justification*. ISA (2006)
7. Kindermann, R., Junttila, T., Niemelä, I.: Modeling for symbolic analysis of safety instrumented systems with clocks. In: *Proc. ACSD 2011*, IEEE (2011) 185–194
8. Manna, Z., Pnueli, A.: *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer (1992)
9. Audemard, G., Cimatti, A., Kornilowicz, A., Sebastiani, R.: Bounded model checking for timed systems. In: *Proc. FORTE 2002*. Volume 2529 of LNCS., Springer (2002) 243–259
10. Sorea, M.: Bounded model checking for timed automata. *Electronic Notes in Theoretical Computer Science* **68**(5) (2002) 116–134
11. Sheeran, M., Singh, S., Stålmarck, G.: Checking safety properties using induction and a SAT-solver. In: *Proc. FMCAD 2000*. Volume 1954 of LNCS., Springer (2000) 108–125
12. Eén, N., Sörensson, N.: Temporal induction by incremental SAT solving. *Electronic Notes in Theoretical Computer Science* **89**(4) (2003) 543–560
13. Bradley, A.R.: SAT-based model checking without unrolling. In: *Proc. VMCAI 2011*. Volume 6538 of LNCS., Springer (2011) 70–87
14. Barrett, C., Sebastiani, R., Seshia, S.A., Tinelli, C.: Satisfiability modulo theories. In: *Handbook of Satisfiability*. IOS Press (2009) 825–885
15. de Moura, L.M., Rueß, H., Sorea, M.: Bounded model checking and induction: From refutation to verification. In: *Proc. CAV 2003*. Volume 2725 of LNCS., Springer (2003) 14–26
16. Dutertre, B., de Moura, L.M.: A fast linear-arithmetic solver for DPLL(T). In: *Proc. CAV 2006*. Volume 4144 of LNCS., Springer (2006) 81–94
17. Kindermann, R., Junttila, T., Niemelä, I.: Beyond Lassos: Complete Bounded Model Checking for Timed Automata. (2012) Currently under review, available at <http://users.ics.tkk.fi/kindermann/2012-CAV-review/bmc.pdf>.
18. Een, N., Mishchenko, A., Brayton, R.: Efficient implementation of property directed reachability. *Proceedings of IWLS*. IEEE/ACM (2011)
19. Nieuwenhuis, R., Oliveras, A.: DPLL(T) with exhaustive theory propagation and its application to difference logic. In: *Proc. CAV 2005*. Volume 3576 of LNCS., Springer (2005) 321–334
20. Cotton, S., Maler, O.: Fast and flexible difference constraint propagation for DPLL(T). In: *Proc. SAT 2006*. Volume 4121 of LNCS., Springer (2006) 170–183
21. Lahtinen, J., Björkman, K., Valkonen, J., Frits, J., Niemelä, I.: Analysis of an emergency diesel generator control system by compositional model checking. *VTT Working Papers 156*, VTT Technical Research Centre of Finland, Espoo (2010)
22. Cimatti, A., Clarke, E.M., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV 2: An opensource tool for symbolic model checking. In: *Proc. CAV 2002*. Volume 2404 of LNCS., Springer (2002) 359–364